



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : H04L 29/06, 12/18, H04N 7/15		A1	(11) International Publication Number: WO 00/36804 (43) International Publication Date: 22 June 2000 (22.06.00)		
(21) International Application Number: PCT/GB99/03871 (22) International Filing Date: 19 November 1999 (19.11.99)		(81) Designated States: AU, CA, JP, SG, US, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).			
(30) Priority Data: 9826158.9 27 November 1998 (27.11.98) GB		Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>			
(71) Applicant (for all designated States except US): BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).					
(72) Inventors; and (75) Inventors/Applicants (for US only): BELL, Sarah [GB/GB]; 9 Brenda Gautrey Way, Cottenham, Cambridge CB4 8XW (GB). ING, Sarom [GB/GB]; 112 Spring Road, Ipswich, Suffolk IP4 2RR (GB). RUDKIN, Steven [GB/GB]; 52 Corder Road, Ipswich, Suffolk IP4 2XD (GB).					
(74) Agent: SHELLEY, Mark, Raymond; BT Group Legal Services, Intellectual Property Dept., Holborn Centre, 8th floor, 120 Holborn, London EC1N 2TE (GB).					
(54) Title: ANNOUNCED SESSION DESCRIPTION					
(57) Abstract					
<p>The invention provides a method of announcing a description of a media session, for example a multimedia conference. In one respect, the invention provides a modular method of announcing media sessions. This method comprises the steps of generating a first base module (410) having a first data structure comprising user oriented data relevant to the media session; generating at least one media module (421, 422, 423) having a second data structure comprising media oriented data necessary for a user to receive a respective media stream of the media session; providing a link between the first base module and the at least one media module; and, announcing the media session by making at least the first base module available to potential recipients of the media session, wherein the link between the first base module and the at least one media module permits a user to access the at least one media module and subsequently receive the media stream.</p>					

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

ANNOUNCED SESSION DESCRIPTION

The present invention relates to the announcement of media stream connections for a media session over a communications network.

5 Multicast transmissions are becoming increasingly common on the Internet. In contrast to standard Internet Protocol (IP) point to point transmissions (unicast), IP multicast allows the simultaneous transmission of information to a group of recipients from a single source. Routing support for IP multicast transmissions is provided by the MBone (IP Multicast Backbone) which is a virtual network layered on top of the
10 Internet.

IP multicast allows real-time communications over wide area IP networks and typical transmissions include video and audio conferencing, live multimedia training, university lectures and transmission of live television programmes.

A multicast transmission usually consists of a multimedia session made up of
15 several individual media streams typically carrying video, audio, whiteboard or raw data. Some sessions are persistent, but the majority exist for a specific period of time, although need not be continuous. Multicast based transmissions on the MBone differ from unicast IP transmissions in that any user receiving the transmission can join the session (unless the transmission is encrypted) and to receive a transmission, a user
20 need only know the appropriate transmission address and timing information.

Prior to a multicast transmission an appropriate announcement containing a session description is made, usually at an IP group multi-cast address. Standard session descriptions are generated using a Session Description Protocol (SDP), as defined in the Internet Engineering Task Force's draft RFC 2327. SDP is a simple ASCII text based
25 protocol that is used to describe real time multimedia sessions and their related scheduling information. SDP messages are wrapped in a carrier protocol, known as a Session Announcement Protocol (SAP), which, in addition to containing the necessary

IP addressing and routing information for transmission across the Internet or MBone, allows the SDP message to be encrypted, signed or compressed. An announcement can then be sent at regular intervals to the announcement group address. As an alternative to SAP, a session may be announced by placing an SDP message on a World Wide Web site (WWW) or by sending it to individuals by email or as a unicast transmission inviting them to participate.

An SDP message conveys information about each media stream in the multicast multimedia session to allow the recipients to participate in the session. A typical SDP message will include the session name and purpose, the time(s) and date(s) the session will be active, the component media streams of the session and information required to participate in each media stream (IP multicast address, port, media format). The SDP message may also include details of the session's bandwidth requirements, an encryption key necessary to participate in a secure multicast transmission using public key encryption, contact information for the organiser of the multicast session, and a Unique Resource Indicator (URI) pointing to a WWW or an Intranet web site where further information on the session may be found, for example, background information relating to the conference.

The level of participation a user may make in a session or stream depends on its purpose. In a multicast television session, typically users would only be able to receive the session streams whilst in a multicast conference session the communication would be bi-directional with a central server (such as group address 120) receiving each participants transmissions and relaying them to the other participants. The level of participation expected of a user in a session or stream may be explicitly stated in the session description or it may be inherent from the session description, for example when a receive-only application is associated with a media stream type in the session description.

A common front end interface used by multicast end users is known as Session Directory Rendezvous (SDR). This interface takes the received announcements, decodes the SDP message and displays the names of those sessions that are still current in a list. The end user may then select one of the listed announcements to view further

5 technical and user-oriented details of the announced session. From the displayed information, the end user can then select to join individual streams of the session or to join the entire session. Once the streams to be joined are selected, SDR starts the necessary multicast-enabled multimedia application on the end user's computer, such as Vic and Vat, and passes the relevant stream information (a transport port address)

10 from the announcement onto the application allowing the application to establish the link to the associated IP multicast address and participate in the stream at transmission time. Having initiated the applications and passed the relevant transport port address SDR plays no further part in the session.

Recent increased usage and demand for (multi)media sessions has highlighted a

15 number of limitations in SDP. SDP limits session descriptions to defining a session having a single set of timings that apply to all of the streams within it. A session in which a stream starts mid-way through the transmission cannot easily be described using SDP. The structure of a session description written in SDP must be a simple linear list of streams which may not reflect the intuitive structure of a complex session. SDP

20 supports a limited and predefined set of applications which can receive the streams and a limited and predefined set of transport mechanisms (e.g. Simple layering, RTP and UDP). As guaranteed Quality of Service (QoS) is becoming more and more desirable to the consumer and the supplier, the need to define QoS policies for the entire session and individual streams in terms of required system resources, bandwidth requirements

25 and supported applications also needs to be met. There may also be requirements on the prioritisation of streams and subsessions or more complicated rules about receiving

streams. A further requirement on the part of the supplier will be the need for charging facilities permitting the charging of an end user for a multicast transmission to which they subscribe according to the QoS and types of streams received etc. There is little scope to include information about QoS policies or charging within the conventional 5 structure of an SDP session description, or any metadata about the session.

According to a first aspect of the present invention there is provided, a method of announcing a description of a media session, comprising the steps of:

generating a first base module having a first data structure comprising user oriented data relevant to the media session;
10 generating at least one media module having a second data structure comprising media oriented data necessary for a user to receive a respective media stream of the media session;

providing a link between the first base module and the at least one media module; and,

15 announcing the media session by making at least the first base module available to potential recipients of the media session,

wherein the link between the first base module and the at least one media module permits a user to access the at least one media module and subsequently receive the media stream.

20 The present invention provides a modular description system for a media session in which session descriptions are constructed in a hierarchical manner providing a plurality of levels of information concerning the constituent parts of the described session.

A problem faced with the current distribution of announcements from the single 25 announcement group address is that there is a limit to the size of each announcement and the frequency with which each can be sent out. In the present invention, it is

possible to provide a modular description system in which a distributed announcement contains links available to the end user to other portions of the announcement which have not been transmitted.

5 Preferably, the method further comprises the steps of: generating a second base module, the second base module containing user orientated data relating to a sub-session of the media session; linking the second base module to the first base module; and, linking said at least one media module to the second base module.

In preferred embodiments, the method further comprises the steps of: generating 10 at least one options module having a third data structure comprising data relating to service level criteria required to participate in the media session; and, linking the or each options module to a respective base module.

The data contained in the options module may relate to a quality of service policy to be used by the media session or a part thereof. Alternatively, the data 15 contained in the options module may relate to a security system to be used by the media session or a part thereof. The data contained in the options module may further relate to a charging system to be used by the media session or a part thereof.

In preferred embodiments, one or more media module(s) comprise data necessary for a user to receive a layered media stream of a respective media session; 20 and said method further comprises the step of linking the or each media module to one or more respective options module(s) containing data relating to a layered mechanism of the respective layered media stream necessary for a party to participate in the layered media stream.

The media session may be announced by transmitting all of the constituent modules of 25 the session description. Alternatively, the media session may be announced by transmitting only some of the constituent modules of the session description, with the

remaining modules of the session description being subsequently accessible by a user using one or more links provided in the modules transmitted. The remaining modules of the session description may be held on one or more servers and the one or more links to the remaining modules are in the form of URI pointers. Modules of the session 5 description contain links to modules which are generated subsequent to the announcement.

According to a second aspect of the invention there is provided a computer readable storage medium containing data defining at least a part of a description of a media session, the session description comprising:-

10 a first base module having a first data structure comprising user oriented data relevant to the media session;

at least one media module having a second data structure comprising media oriented data necessary for a user to receive a respective media stream of the media session;

15 a link between the first base module and the at least one media module;

Another problem faced by providers of current (multi)media sessions and the developers of the associated (multi)media applications is the spread of skills required to implement an application that can initiate and manage a real-time data connection over a communications network and perform the (multi)media functions the end user would 20 expect. For example, developers of multimedia applications require teams with skills in audio and video coding, network transport protocols, real time programming, user interface design and integration techniques. The session description of the present invention simplifies this process by allowing the necessary communication channels and media streams to be identified in the session description. This information is used by 25 generic middleware in the form of a session control and communications manager to dynamically instantiate the respective streams and channels for the applications at run

time.

Furthermore, until now the only way a QoS policy could be implemented was to process a session description to determine which streams of a session could or should be run and then to initialise the applications so they connect to the respective streams.

5 This required the communications manager not only to know about the session requirements and available system resources but also the capabilities of each application.

In a preferred example of the present invention the media modules of a session description are checked by the respective multimedia client application prior to QoS management, thereby reducing the workload of the communications manager, that is to say the respective client applications determine whether the media modules can be supported. Furthermore, applications need only request streams from the session control system associated with the client since the session control now handles centrally the creation and management of streams in real time. This aspect is also the

15 subject of our co-pending UK patent application 9826157.1.

The present invention simplifies application development and service provision. A further problem is that applications should be able to adapt to available network and host resources. This is particularly important for multi-party applications operating in heterogeneous environments where each party may have different resources available to them. Furthermore the nature of the heterogeneity may vary over the lifetime of the session, for example as network congestion varies or as the terminal resources are shared with other applications or other users. The present invention is able to use a QoS policy incorporated within the session description to prioritise the allocation of resources and to determine whether participation in the session is viable.

20

25 A still further problem is that the application developer and service provider typically need to address security and charging requirements. The present invention

allows security and charging policies to be incorporated within the session description for use within the session control system to invoke appropriate charging and security procedures. Instead of having to develop security and charging functions the application developer and service provider need only specify appropriate policies.

5 In the present invention application development is simplified by using the session description to drive the dynamic management of communication channels and to adapt to available resources. It also reduces the problem of handling charging and security requirements to a matter of specifying charging and security policies within the session description.

10 An example of the present invention will now be described in detail with reference to the accompanying drawings, in which:

Figure 1 is a schematic diagram illustrating a multicast transmission across the MBone;

15 Figure 2 is a schematic diagram illustrating the distribution of an SDP announcement;

Figure 3 is a block diagram of a modular session description of a simple session generated in accordance with the present invention;

Figure 4 is a block diagram of a modular session description of a complex session generated in accordance with the present invention;

20 Figure 5 is a schematic diagram of a system for managing media stream connections;

Figure 6 is a flow chart illustrating the steps involved in managing a media session according to the system of Figure 5; and,

Figure 7 is a flow chart further illustrating a parsing step of Figure 6.

25 An example of an IP multicast transmission system is described with reference to Figure 1. Prior to a multicast transmission, an appropriate announcement containing

a session description is made, thereby allowing end users 110a-110e to elect to receive the transmission. Each end user electing to receive the transmission is linked to a group IP Multicast address 120 associated with the transmission. At the transmission time of the multicast session, the session streams are transmitted from a source 130, or a plurality of sources, to the group address. At the group address, the transmission is disseminated along the links 140 to those end users who have elected to receive it (in this example end users 110a-110c).

An example of an announcement and election system is described with reference to Figure 2. Most public multicast sessions are announced at a single group IP 10 multicast address 200 dedicated to the transmission of announcements to multicast end users. End users 210a-210e electing to receive the announcements are linked to the announcement group address and, in the same way as an actual session transmission, each announcement arriving at the announcement group address is disseminated to the end users. A front end interface 220 on each end user's computer 15 displays information obtained from the associated session description for each announcement. The minimum information a session description may contain is a time and date that the session will be active and the group IP multicast address(es) from which the end user may elect to receive one or more media streams and to which they could send their own streams for the session. Using the front end interface, an end user 20 can select the announced session(s), or their component stream(s) they wish to participate in.

Figure 3 is a block diagram of a session description 300 for a simple multicast television session. The session description 300 comprises a base module 310 linked to a media module 320.

25 The base module 310 contains user oriented data relating to the session including the title and timing information. The base module 310 may also include a

description or abstract, contact information about the organiser and a WWW or an intranet URI pointing to a web site containing further information. Ideally, the base module 310 should contain enough information for the user to decide if they are interested in participating in the session.

5 The media module 320 contains announcement data relating to a video stream of the session. The media module 320 contains the technical information (data) necessary for the user to receive the associated media stream. In particular, connection, timing and media format details are provided.

A first example of a session description 300 generated for transmission to end
10 users is shown below:

```
(  
  type=(base)  
  id=(310)  
 15  info=(title="live multicast television session")  
    source=(name="A.Sender" email=asender@tx.com)  
    media=(video=(client=odbits0.16))  
    time=(length=50m repeat=continuous)  
    category=("Entertainment")  
 20  options=(none)  
    modules=(m=320)  
  )  
  
(  
  type=(media)  
  id=(320 310)  
  media=(video=(client=odbits0.16))  
  connection=(229.1.1.2/7000)  
 30  time=(length=50m)  
 )
```

Session description example 1

The base module 310 has a unique identifier (id field) used in the generation of
35 links between two modules during the processing of the session description. The
modules field of the base module 310 lists the type and unique identifier of the media

module 320 linked to the base module 310. The second identifier in the id field of the media module 320 is the unique identifier belonging to the base module 310 linking the media module back to the base module 310. By extension, these two-way links permit a module tree to be traversed from a base module downwards or from a media module upwards. The use of this feature is described later with reference to session description example 4.

The connection field of the media module 320 contains the IP multicast address and port number from which the media stream can be received.

Figure 4 is a block diagram of a session description 400 for a complex multicast session of a multimedia conference with two tracks, or sub-sessions, and a panel discussion. Each track provides multiparty video and audio conferencing and a shared whiteboard for leaving notes and messages. The panel discussion is encrypted and the whole conference is subject to a subscription fee payable in advance by each participant.

The session description 400 contains a top level base module 410 linked to further base modules 420, 430, 440 and an options module 411. The top level base module 410 contains data relating to the overall session including its name, purpose and timing information. The options module 411 contains details of the payment mechanism for subscription fees.

Each further base module 420, 430, 440 relates to a subsession of the conference. Base module 420 relates to the first track of the conference. The base module 420 is linked to media modules 421-423, each containing connection, timing and media format data for respective video, audio and whiteboard streams.

The base module 420 is also linked to options module 424 which contains data relating to a QoS policy for the first track defining which media modules are optional and which are mandatory for a participant of the first track. The mandatory list contains

identifiers of those media modules which are needed for the session or subsession to operate correctly whilst the optional list contains identifiers of the media modules that are not necessary for the session or subsession to operate correctly if system resources are scarce.

5 The base module 430 relates to the second track of the conference. It is linked to media modules 431-433, each containing connection, timing and media format details for respective video, audio and whiteboard streams. The base module 430 is also linked to options module 434 which contains data relating to a QoS policy for the second track defining which media modules are optional and which are mandatory for a
10 participant of the second track. Base module 440 relates to the panel discussion. It is linked to media modules 441 and 442, each containing connection, timing and media format details for respective video and audio streams of the panel discussion. The base module 440 is also linked to options module 443 which contains encryption details (ie. how and where to get the necessary cryptographic keys) necessary for a participant to
15 decode the panel discussion media streams 441, 442 according to a known encryption mechanism such as DES or public key encryption.

 The video media stream defined in media module 441 is layered. Layering of media streams allows users with different system resources to receive as much of the stream as their system resources allows. Every user must receive the bottom layer of
20 the stream containing the minimum stream data. However, if a user has sufficient free system resources they can receive the next layer up containing enhancements to the previous layer. Successive layers can be received enhancing the received media stream until the maximum number of layers is received or all free system resources capacity is used. The media module 441 is linked to an options module 444 which contains data
25 on the layering necessary for the end user to be able to receive the layered stream correctly.

The portion of the session description 400 generated for modules 410, 411, 420 and 440 for transmission to end users is shown below in session description example 2.

```
5      ( # overall conference session
  type=(base)
  id=(410)
  info=(title="Multimedia98 Conference")
  source=(owner="Joe Bloggs" email=joe@nowhere.com)
10     media=(video=(client=RealPlayerG2) whiteboard=(client=wb))
  time(start="09:00 GMT 25/12/98" stop="13:00 GMT 25/12/98")
  options=(oc=411)
  modules=(b=420 b=430 b=440 oc=411)
)
15     (
  # conference track 1
  type=(base)
  id=(420 410)
  info=(title="MM98 Systems and Applications Track")
  source=(owner="Joe Bloggs" email=joe@nowhere.com)
20     media=(video=(client=RealPlayerG2) whiteboard=(client=wb))
  time(start="09:00 GMT 25/12/98" stop="11:00 GMT 25/12/98")
  options=(osq=424)
  modules=(m=421 m=422 m=423 osq=424)
)
25     (
  # session QoS for track 1
  type=(option-sQoS)
  id=(424 420)
  mandatory=(421 422)
  optional=(423)
)
30     (
  # conference panel discussion
  type=(base)
  id=(440 410)
  info=(title="MM98 Panel Discussion")
  source=(name="Joe Bloggs" email=joe@nowhere.com)
40     media=(video=(client=RealPlayerG2) whiteboard=(client=wb))
  time(start="11:00 GMT 25/12/98" stop="13:00 GMT 25/12/98")
  options=(osec=443)
  modules=(m=441 m=442 osec=443)
)
45     (
  # video for panel discussion
  type=(media)
  id=(441 440)
  info=(title="MM98 Panel Discussion Video")
  source=(owner="Joe Bloggs" email=joe@nowhere.com)
50     media=(video=(type=live client=RealPlayerG2))
  connection=(226.0.0.106/1010 policy=444)
```

```
time=(start="11:00 GMT 25/12/98" stop="13:00 GMT 25/12/98")
)
(
  # media QoS policy for panel discussion video
5   type=(option-mQoS)
    id=(444 440)
    mechanism=(layer=(base=226.0.0.106/1010 number=3))
)
10  (
    # encryption policy for panel discussion
    type=(option-sec)
    id=(443 440)
    participant=(member=w3c)
    publickey=(location=http://www.w3.org/members_only/)
15  info=(location=http://www.w3.org/)
)
(
  # charging policy for entire conference
  type=(option-chg)
20  id=(411 410)
    mechanism=(type=AAA)
    price=(fee=1000GBP)
    info=(location=http://www.aaa.net/)
)
25  Session description example 2
```

Where there is surplus network bandwidth available, complete session descriptions can be announced to end users who may then elect to receive the announced session or parts thereof. However, the individual modules of the session 30 description do not need to be announced together. If the network bandwidth available for announcements restricts the size of session descriptions, only the top level base module may be announced. In this situation, the link between modules may be, for example, a URI to a WWW or an intranet web site or server, an email address, an IP multicast address, an FTP address or details of a file or database stored on a local 35 computer system from which an interested user can obtain the remaining modules.

The following session description example illustrates how the above session description for base module 420 would be changed if media module 421 was stored on a WWW server:

```
5      ( # conference track 1
        type=(base)
        id=(420 410)
        info=(title="MM98 Systems and Applications Track")
        source=(owner="Joe Bloggs" email=joe@nowhere.com)
        media=(video=(client=RealPlayerG2) whiteboard=(client=wb))
        time(start="09:00 GMT 25/12/98" stop="11:00 GMT 25/12/98")
        options=(osq=424)
        modules=(m=421 location=http://www.announce.org/cgi-bin/module.cgi?id=421
10      m=421 m=423 osq=424)
      )
      Session description example 3
```

Furthermore, top level modules of a session description may be announced well 15 in advance of the actual transmission, at a time where the final details of content are unknown, in which case the remaining levels may be made available from pre-announced links at a later time.

Figure 5 is a schematic diagram of a system for managing media stream connections at a terminal of an end user system according to the present invention.

20 The session control system 500 is linked to an announcement receiving interface 510 and one or more multicast-capable multimedia applications 520. The session control system 500 and the announcement receiving interface 510 are connected to a network interface 530 via which announcements may be received and multicast transmissions may be initiated and/or received.

25 Announcements received at the network interface 530 are routed to the receiving interface 510. The receiving interface 510 decodes each announcement to obtain the session description and displays the user oriented information from the one or more base modules in a list to the user. The user is able to select a session description from the list announcing a session they wish to receive. The selected 30 description is passed to the session control system 500 which determines which of the user's multimedia applications 520 are required for participation in the described session, starts the applications and initiates and provides the necessary media streams

to the respective applications 520 via a communications manager 550.

The receiving interface 510 may be linked to other Internet communications applications 540 such as a WWW browser or an email client (not shown) which may be used to gather further information on the described session based on links provided in 5 the session description. Also, where an incomplete set of base and/or media modules of a session description are received, the receiving interface 510 attempts to obtain the remaining modules using the Internet communications applications prior to passing it onto the session control system 500.

Figure 6 is a flow chart showing the steps taken by the session control system 10 500 upon receipt of a session description. The description is first parsed in step 600 to identify client applications for each media module. Once this is done a second parse is carried out where applications are launched in step 610, that is to say for each media module start the application specified in the client field if that application has not already been started. The portion of the session description relating to the respective 15 media type, i.e. the media module, the base module directly above the media module, all other modules attached to that base module and any other options modules that apply, is passed to the corresponding application in step 620. Since the media modules are marked with appropriate client applications, each application will be able to select those media streams that it wants to participate in. The application replies to the 20 session control system with a connection request specifying its requirements in the form of a list of identifiers of media modules from which streams are to be initiated in step 630. The connection request is assembled by the session control system in step 640 and the system then parses the session description to identify other applications to launch in step 645. If a further media type is found, steps 610 to 640 are repeated, 25 otherwise the session control system uses the assembled connection requests to form a list of media modules. This list is passed, together with a session QoS policy, to the

communications manager, a system used in by the session control system, which determines according to the QoS policies and available system resources whether each connection request is viable.

The session QoS policy is constructed in two steps:- first, the multiple session 5 QoS policies relevant for all the media modules to be initiated are combined into one session QoS policy: second, the resulting session QoS policy may be adapted to take account of (a) user default preferences (defined in a user profile), (b) a user's wish to determine the policy interactively, and (c) an application's default configuration (defined in the application profile(s)).

10 The communications manager responds to the session control system in step 650 with an indication of the viable media stream connection requests. If necessary, the session control system may contact a charging system to initiate accounting for the session prior to requesting the communications manager to create the viable media stream connections in step 660.

15 Once a session starts, each received data stream relating to the session is passed to the associated multimedia application in step 670 until the scheduled stream time ends in step 680 or the multimedia application requests to the session control system that the connection is terminated in step 690, at which point the session control system disconnects the connection in step 700.

20 Figure 7 is a flow chart showing the QoS management step 650 of Figure 6 in greater detail.

Having received the assembled list of connection requests, the communications manager matches each item of this list to a media profile in step 705. A media profile defines requirements which must be met for the requested media stream to operate on 25 the end user's computer including the minimum network bandwidth needed for satisfactory reception of the stream.

A terminal profile is determined in step 710. The terminal profile defines the resources which are available at the end user's computer for use by the requested media streams. This includes available network bandwidth, free memory and disk space and available hardware such as monitor size, processor speed and free audio and video capture devices. The media profile of each connection request is compared against the available system resources defined by the terminal profile in step 720. If the terminal profile matches or exceeds the media profile, the connection request is declared viable in step 730 and the terminal profile is decremented accordingly for the remaining connection requests in step 740. Each connection request is processed until there are no remaining requests or until the media profile of a request exceeds the terminal profile. In this situation, the communications manager determines the optimum terminal profile the user's computer would have if all non-essential applications were not running in step 750 and whether the computer is capable of fulfilling the media profile in step 760. If the computer is capable of fulfilling the media profile, the communications manager attempts to free system resources from currently allocated streams or connection requests which have lower priority or by asking the user to terminate other non-essential applications running on the computer in step 770. Alternatively, this could be done by reducing the number of layers received from a layered stream transmission. If sufficient resources cannot be found an exception is reported to the user and the connection request is marked as unviable. If the media stream that cannot be received is defined as mandatory in a QoS policy for a media session or subsession, all the connection requests for that media session or subsession are cancelled in step 790. If, however, the media stream is optional, the communications manager continues processing further connection requests in step 720. Once all pending connection requests have been processed, the communications manager reports those that are viable to the session control system.

The processing of a session description will now be described with reference to Figure 4 and session description example 4 which is the session description generated for Track 1 (modules 410 and 420-424 of Figure 4).

```

5      ( # overall conference session
10     type=(base)
      id=(410)
      info=(title="Multimedia98 Conference")
      source=(owner="Joe Bloggs" email=joe@nowhere.com)
      media=(video=(client=RealPlayerG2) whiteboard=(client=wb))
      time(start="09:00 GMT 25/12/98" stop="13:00 GMT 25/12/98")
      options=(oc=0010)
      modules=(b=420 b=430 b=440 oc=411)
    )
15
    ( # conference track 1
      type=(base)
      id=(420 410)
      info=(title="MM98 Systems and Applications Track")
      source=(owner="Joe Bloggs" email=joe@nowhere.com)
      media=(video=(client=RealPlayerG2) whiteboard=(client=wb))
      time(start="09:00 GMT 25/12/98" stop="11:00 GMT 25/12/98")
      options=(osq=424)
      modules=(m=421 m=422 m=423 osq=424)
    )
25
    ( # video for track 1
      type=(media)
      id=(421 420)
      info=(title="MM98 Systems and Applications Track Video")
      source=(owner="Joe Bloggs" email=joe@nowhere.com)
      media=(video=(type=live client=RealPlayerG2))
      connection=(226.0.0.100/1000)
      time=(start="09:00 GMT 25/12/98" stop="11:00 GMT 25/12/98")
    )
30
    ( # audio for track 1
      type=(media)
      id=(422 420)
      info=(title="MM98 Systems and Applications Track Audio")
      source=(owner="Joe Bloggs" email=joe@nowhere.com)
      media=(audio=(type=live format=g711))
      connection=(226.0.0.101/1001)
      time=(start="09:00 GMT 25/12/98" stop="11:00 GMT 25/12/98")
    )
40
    ( # whiteboard for track 1
      type=(media)
      id=(423 420)
      info=(title="MM98 Systems and Applications Track Whiteboard")
      source=(owner="Joe Bloggs" email=joe@nowhere.com)
    )
45
50

```

```
media=(whiteboard=(client=wb))
connection=(226.0.0.102/1002)
time=(start="09:00 GMT 25/12/98" stop="11:00 GMT 25/12/98")
)
5
(
# session QoS for track 1
type=(option-sQoS)
id=(424 420)
mandatory=(421 422)
10
optional=(423)
)
```

Session description example 4

15 The session control system, having received the above session description, processes the tree structure of the session description starting at base module 410. The first module encountered is base module 420. As this is not a media module but it does have sub-modules, the session control system continues down this branch to media module.

20 The media field of the media module 421 already defines the multimedia client application required as RealPlayerG2 (a multimedia application of Real Networks Inc) thus the session control system ignores it and continues to the next media module. The media field of the media module 422 does not have a multimedia client application defined, however a format for the audio data is specified. The session control system 25 recognises that this particular audio format can be supported by RealPlayerG2 so it amends the media field to read client=RealPlayerG2. The next media module 423 has already defined a client application as wb so it ignores this module, and it also ignores the option module 424.

30 The session control system parses the tree structure again in order to launch client applications. The first media module 421 specifies that RealPlayerG2 should be launched, hence the session control system launches the application on the end user's system and keeps a record of this activity. The second media module 422 specifies an

application that has already been launched and so the session control system ignores it and continues to the next media module. The media module 423 specifies that wb should be launched, so the session control system launches the application and keeps a record of this activity.

5 RealPlayerG2 is passed the media module 421, base module 420 and modules 422-424. The application processes the media modules given to determine which it can handle, and in this case it identifies 421 and 422. Having determined which streams it can handle, the application sends a connection request back to the session control system requesting connection to the media streams of modules 421 and 422. Similarly, 10 wb is passed the media module 423, base module 420, modules 421-422, and the module 424. The application processes the given modules as described previously, and requests connection to the media stream of modules 423.

The above connection requests are assembled by the session control system into a list, this list is then passed to the communications manager along with the 15 session QoS policy module 424. The communications manager determines whether each request is viable according to the steps of Figure 7.

Assuming there are sufficient resources for all the connection requests for mandatory media streams, the communications manager passes back a list of viable streams to the session control system which then processes the tree again to determine 20 the connection data held in the connection field of each media module so it can request that the communications manager initiate a connection to the appropriate media stream for each of the viable connection requests according to the connection data. The session control system then manages the session and its media stream connections as is described with reference to steps 670 to 700 of Figure 6.

25 Due to the heterogeneity of the Internet and differing capabilities and operating environments of end user computer systems, the session control system described has

been implemented in Java (Java is a Trade Mark of Sun Microsystems Inc.). The announcement receiving interface, Session Directory, receives the announcements and passes those selected by the end user to the session control manager implemented as an application programming interface running as a background process on the end 5 user's computer.

Whilst the present invention has been described with reference to the Internet and multicast transmissions, it will be apparent to the reader that the described modular session description and the session control system are applicable to the announcement and subsequent management of connections to media streams of a (multi)media session 10 using other known transport mechanisms such as unicast.

Furthermore, although mechanisms for encryption, charging and other such services have not been explicitly described, it would be apparent to the reader that appropriate session descriptions and associated functions within the session control system for their processing could be readily implemented according to the mechanism 15 required.

CLAIMS

1. A method of announcing a description of a media session, comprising the steps of:

5 generating a first base module having a first data structure comprising user oriented data relevant to the media session;

generating at least one media module having a second data structure comprising media oriented data necessary for a user to receive a respective media stream of the media session;

10 providing a link between the first base module and the at least one media module; and,

announcing the media session by making at least the first base module available to potential recipients of the media session,

wherein the link between the first base module and the at least one media 15 module permits a user to access the at least one media module and subsequently receive the media stream.

2. A method according to claim 1, further comprising the steps of:

generating a second base module, the second base module containing user 20 orientated data relating to a sub-session of the media session;

linking the second base module to the first base module; and,

linking said at least one media module to the second base module.

3. A method according to claim 1 or 2, further comprising the steps of:

25 generating at least one options module having a third data structure comprising data relating to service level criteria required to participate in the media session; and,

linking the or each options module to a respective base module.

4. A method according to claim 3, in which the data contained in the options module relates to a quality of service policy to be used by the media session or a part 5 thereof.

5. A method according to claim 3 or 4, in which the data contained in the options module relates to a security system to be used by the media session or a part thereof.

10 6. A method according to any of claims 3 to 5, in which the data contained in the options module relates to a charging system to be used by the media session or a part thereof.

7. A method according to any preceding claim, wherein one or more media 15 module(s) comprise data necessary for a user to receive a layered media stream of a respective media session; and said method further comprises the step of linking the or each media module to one or more respective options module(s) containing data relating to a layered mechanism of the respective layered media stream necessary for a party to participate in the layered media stream.

20

8. A method according to any preceding claim, in which the data contained in a media module includes data necessary for a user to receive or transmit data or both receive and transmit for inclusion in the media session.

25 9. A method according to any preceding claim, in which the media session is announced by transmitting all of the constituent modules of the session description.

10. A method according to any of claims 1 to 8, in which the media session is announced by transmitting only some of the constituent modules of the session description, with the remaining modules of the session description being subsequently accessible by a user using one or more links provided in the modules transmitted.

11. A method according to claim 10, in which the remaining modules of the session description are held on one or more servers and the one or more links to the remaining modules are in the form of URI pointers.

10

12. A method according to any preceding claim, in which modules of the session description contain links to modules which are generated subsequent to the announcement.


15 13. A computer readable storage medium containing data defining at least a part of a description of a media session, the session description comprising:-

16 a first base module having a first data structure comprising user oriented data relevant to the media session;

17 at least one media module having a second data structure comprising media oriented data necessary for a user to receive a respective media stream of the media session;

18 a link between the first base module and the at least one media module; wherein the link permits a user to access the at least one media module and subsequently receive the media stream.

Fig. 1.



2/7

Fig.2.





Fig.3.



3/7

Fig.4.






Fig.5.




5/7




6/7

Fig.7.



7/7

Fig.7 (Cont.).



INTERNATIONAL SEARCH REPORT

Interr. Application No
PCT/GB 99/03871

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L29/06 H04L12/18 H04N7/15

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L H04N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	AVARO O ET AL: "The MPEG-4 systems and description languages: A way ahead in audio visual information representation" SIGNAL PROCESSING. IMAGE COMMUNICATION, NL, ELSEVIER SCIENCE PUBLISHERS, AMSTERDAM, vol. 9, no. 4, 1 May 1997 (1997-05-01), pages 385-431, XP004075337 ISSN: 0923-5965 abstract page 387, line 17 -page 393, line 2 page 399, line 22-43 page 415, line 6 -page 417, line 7 figure 5 ---- -/-	1,2,8,9, 13
A		3,7,10, 11

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

& document member of the same patent family

Date of the actual completion of the international search

20 April 2000

Date of mailing of the international search report

03/05/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Lievens, K

INTERNATIONAL SEARCH REPORT

Intern	ai Application No
PCT/GB 99/03871	

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 798 905 A (DIGITAL VISION LAB CORP) 1 October 1997 (1997-10-01) abstract column 1, line 24 -column 2, line 48 column 4, line 12 -column 6, line 24 column 8, line 49 -column 9, line 5 column 10, line 55 -column 11, line 15 figures 3-5,15 -----	1,8,13
X	THIMM H ET AL: "A MAIL-BASED TELESERVICE ARCHITECTURE FOR ARCHIVING AND RETRIEVING DYNAMICALLY COMPOSABLE MULTIMEDIA DOCUMENTS" MULTIMEDIA TRANSPORT AND TELESERVICES. INTERNATIONAL COST 237 WORKS PROCEEDINGS, VIENNA, NOV. 13 - 15, 1994, 13 November 1994 (1994-11-13), pages 14-34, XP000585292 HUTCHISON D;DANTHINE A; LEOPOLD H; COULSON G (EDS) abstract page 16, line 20 -page 19, line 9 -----	1,2,8,13
X	"SDP: Session Description Protocol" RFC2327, April 1998 (1998-04), pages 1-42, XP002101463 http://www.cis.ohio-state.edu/htbin/rfc/rfc2327.html cited in the application page 6, line 20 -page 8, line 6 page 17, line 15 -page 18, line 38 -----	1,9,13
A	WO 97 22201 A (XIE DONG ;CAMPBELL ROY H (US); CHEN ZHIGANG (US); TAN SEE MONG (US) 19 June 1997 (1997-06-19) abstract page 9, line 10-15 page 12, line 18 -page 20, line 4 page 36, line 17 -page 37, line 11 -----	21,39, 59,60
A		1,13

INTERNATIONAL SEARCH REPORT**Information on patent family members**

Interr. nat Application No

PCT/GB 99/03871

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
EP 0798905	A 01-10-1997	JP	9247142 A	19-09-1997
		CA	2199103 A	05-09-1997
		CN	1170164 A	14-01-1998
		US	5758086 A	26-05-1998
WO 9722201	A 19-06-1997	EP	0867003 A	30-09-1998